

Automatic numerical differentiation by maximum likelihood estimation of state-space model [★]

Robert Piché ^a

^a*Tampere University of Technology, Tampere, Finland*

Abstract

A linear Gaussian state-space smoothing algorithm is presented for estimation of derivatives from a sequence of noisy measurements. The algorithm uses numerically stable square-root formulas, can handle simultaneous independent measurements and non-equally spaced abscissas, and can compute state estimates at points between the data abscissas. The state space model's parameters, including driving noise intensity, measurement variance, and initial state, are determined from the given data sequence using maximum likelihood estimation computed using a expectation maximisation iteration. In tests with synthetic biomechanics data, the algorithm has equivalent or better accuracy compared to other automatic numerical differentiation algorithms.

Key words: Smoothing filters, estimation algorithms, data analysis, regularization

1 Introduction

Numerical differentiation (ND) of a sequence of noisy measurements is an important problem in data analysis. For example, one may want to estimate velocity and acceleration from a sequence of displacement measurements. The problem has been well studied; comparative surveys of ND algorithms include [11,3,5,4,8,15,1,13].

Because differentiation amplifies noise, catastrophically so when the sampling rate is high, an effective ND method must trade off data fidelity with noise smoothing. In most ND algorithms, the trade-off is governed by one or more user-defined parameters, variously called regularisation, smoothing, or bandwidth (cutoff frequency) parameters. Some ND algorithms are “automatic”, in the sense that they determine the smoothing parameters for a given time series without knowledge of the true signal values. The surveys [4,8] assess several automatic ND algorithms.

Numerical differentiation can be approached as a standard state space estimation problem with continuous-time dynamics and discrete-time measurements. In the Kalman fixed-lag smoother of Fioretti and Jetto [6,7], the state space dynamic model is a multiply-integrated stationary Wiener process, and the measurement error is an additive stationary discrete-time Gaussian white noise. In the target tracking literature

this family of state-space models is known as the polynomial motion model [2, §6.2], of which the constant velocity model is the best known example.

The ND algorithm presented here is also based on the state space model of the multiply-integrated stationary Wiener process. Derivatives are estimated using fixed-lag Rauch-Tung-Striebel smoothing implemented with numerically stable square-root formulas. The algorithm can treat independent simultaneous measurements and non-equally-spaced abscissas, and supports evaluation at abscissas other than data points (“dense output”). A maximum likelihood (ML) estimate of all the state space models’ parameters, namely the initial state, driving noise intensity, and measurement noise variance, is computed using an extension of the expectation-maximisation (EM) algorithm for state space model identification [14,9]. A MATLAB implementation of the algorithm is freely available for download ¹.

2 Algorithm

2.1 Signal model

The underlying signal is assumed to be the $(d - 1)$ -fold integral of a Wiener process. The linear stochastic differential equation is

$$dx = Fx dt + qLdw$$

[★] This paper was not presented at any IFAC meeting. Corresponding author R. Piché. Tel. +358 40 8490174.

Email address: robert.piche@tut.fi (Robert Piché).

¹ <https://se.mathworks.com/matlabcentral/fileexchange/xxxxxxx>

where w is the standard Wiener process, the underlying signal is the first component of the d -component state vector x , its first derivative is the second component, etc., and

$$F = \begin{bmatrix} 0 & 1 & 0 & 0 & \cdots & 0 \\ 0 & 0 & 1 & 0 & \cdots & 0 \\ \vdots & & & & & \vdots \\ 0 & 0 & 0 & 0 & \cdots & 1 \\ 0 & 0 & 0 & 0 & \cdots & 0 \end{bmatrix}, \quad L = \begin{bmatrix} 0 \\ 0 \\ \vdots \\ 0 \\ 1 \end{bmatrix}.$$

The parameter $q > 0$ is the intensity (spectral density) of the driving white noise.

The abscissas for the discrete-time state space model are denoted t_k for $k = 1, 2, \dots$; the sequence t_1, t_2, \dots is assumed to be monotonically increasing. Denoting $x_k = x(t_k)$ and $\Delta_k = t_{k+1} - t_k$, the discrete-time dynamic model is a linear state space model driven by additive discrete Gaussian white noise

$$x_{k+1}|x_k \sim N(A_k x_k, Q_k), \quad k = 1, 2, \dots, \quad (1)$$

where $N(\cdot, \cdot)$ denotes a Gaussian distribution with given mean and covariance, the dynamic transition matrix is

$$A_k = \exp(F\Delta_k) = \begin{bmatrix} 1 & \Delta_k & \frac{1}{2!}\Delta_k^2 & \cdots & \frac{1}{(d-1)!}\Delta_k^{d-1} \\ 0 & 1 & \Delta_k & \cdots & \frac{1}{(d-2)!}\Delta_k^{d-2} \\ \vdots & & & & \vdots \\ 0 & 0 & 0 & \cdots & 1 \end{bmatrix},$$

and the driving noise covariance is $Q_k = q\bar{Q}_k$, where

$$\begin{aligned} \bar{Q}_k &= \int_0^{\Delta_k} \exp(F(\Delta_k - \tau)) L L^T \exp(F^T(\Delta_k - \tau)) d\tau \\ &= \text{diag}\left(\begin{bmatrix} \Delta_k^{2d-1} \\ \vdots \\ \Delta_k^{1/2} \end{bmatrix}\right) \begin{bmatrix} \frac{1}{2d-1} & \cdots & \frac{1}{4} & \frac{1}{3} \\ \frac{1}{2d-2} & \cdots & \frac{1}{3} & \frac{1}{2} \\ \vdots & & \vdots & \vdots \\ \frac{1}{d} & \cdots & \frac{1}{2} & 1 \end{bmatrix} \text{diag}\left(\begin{bmatrix} \Delta_k^{2d-1} \\ \vdots \\ \Delta_k^{1/2} \end{bmatrix}\right). \end{aligned}$$

At each abscissa, there are n_k scalar measurements, denoted $y_{k,1}, \dots, y_{k,n_k}$. Each measurement is modelled as the signal value plus independent additive zero-mean Gaussian noise, that is,

$$y_{k,j}|x_k \sim N(Hx_k, R), \quad j = 1, \dots, n_k, \quad (2)$$

where $H = [1, 0, \dots, 0]$ and R is the variance.

2.2 Fixed-Interval Smoothing

Let $x_{k|j}$ denote the state conditioned on the measurements at times t_1, \dots, t_j . For the linear Gaussian state space model

(1) and (2), and a Gaussian prior distribution

$$x_{1|0} \sim N(m_{1|0}, P_{1|0}), \quad (3)$$

all posterior states $x_{k|j}$ are jointly Gaussian. *Fixed-interval smoothing* is the computation of the mean and covariance of the states $x_{1|T}, \dots, x_{T|T}$, given the model parameters

$$\theta = [q, R, m_{1|0}, P_{1|0}]$$

and the measurements $y_{1:T} = [y_{1,1}, \dots, y_{T,n_T}]$. The Rauch-Tung-Striebel (RTS) smoother computes these states sequentially, with a forward pass (a Kalman filter) that processes the measurements, followed by a backward pass. For better numerical stability, the QR factorisation-based square root RTS algorithm of [9] is used, as follows.

The *forward pass* consists of two stages that are carried out for each $k = 1, 2, \dots, T$. Before the beginning of the forward pass, $P_{1|0}^{1/2}$, the lower triangular Cholesky factor of $P_{1|0}$, is computed. The first stage, the *measurement update*, is the computation of the parameters of the filtering distribution $x_{k|k} \sim N(m_{k|k}, P_{k|k}^{1/2} P_{k|k}^{T/2})$ by the formulas

$$m_{k,0|k} = m_{k|k-1}, \quad P_{k,0|k}^{1/2} = P_{k|k-1}^{1/2}$$

for $j = 1, \dots, n_k$ do

\mathcal{R} = triangular factor of QR decomposition of

$$\begin{bmatrix} R^{1/2} & H P_{k,j-1|k}^{1/2} \\ 0 & P_{k,j-1|k}^{1/2} \end{bmatrix}^T$$

$$S_{k,j} = \mathcal{R}_{1,1}^T \mathcal{R}_{1,1}$$

$$K_{k,j} = \mathcal{R}_{1,2:d+1}^T \mathcal{R}_{1,1}^{-T}$$

$$P_{k,j|k}^{1/2} = \mathcal{R}_{2:d+1,2:d+1}^T$$

$$v_{k,j} = y_{k,j} - H m_{k,j-1|k}$$

$$m_{k,j|k} = m_{k,j-1|k} + K_{k,j} v_{k,j}$$

end do

$$m_{k|k} = m_{k,n_k|k}, \quad P_{k|k}^{1/2} = P_{k,n_k|k}^{1/2}$$

The second stage, the *dynamic update*, is the computation of the parameters of the one-step prediction distribution $x_{k+1|k} \sim N(m_{k+1|k}, P_{k+1|k}^{1/2} P_{k+1|k}^{T/2})$. The formulas for the dynamic update are

$$m_{k+1|k} = A_k m_{k|k}$$

\mathcal{R} = triangular factor of QR decomposition of

$$\begin{bmatrix} P_{k|k}^{T/2} A_k^T \\ Q_k^{1/2} \end{bmatrix}^T$$

$$P_{k+1|k}^{1/2} = \mathcal{R}_{1:d,1:d}^T$$

This stage is omitted for $k = T$.

In the *backward pass*, the parameters of the joint smoothing distribution

$$\begin{bmatrix} x_{k+1|T} \\ x_{k|T} \end{bmatrix} \sim \mathcal{N}\left(\begin{bmatrix} m_{k+1|T} \\ m_{k|T} \end{bmatrix}, \begin{bmatrix} P_{k+1|T} & P_{k+1|T} G_k^T \\ G_k P_{k+1|T} & P_{k|T} \end{bmatrix}\right)$$

are computed sequentially for $k = T-1, \dots, 1$; the smoothing distribution is then $x_{k|T} \sim \mathcal{N}(m_{k|T}, P_{k|T})$. The backward pass formulas are

$$\begin{aligned} P_{k|k} &= P_{k|k}^{1/2} P_{k|k}^{T/2} \\ P_{k+1|k}^{-1} &= P_{k+1|k}^{-T/2} P_{k+1|k}^{-1/2} \\ G_k &= P_{k|k} A_k^T P_{k+1|k}^{-1} \\ m_{k|T} &= m_{k|k} + G_k(m_{k+1|T} - m_{k+1|k}) \\ \mathcal{R} &= \text{triangular factor of QR decomposition of} \\ &\quad \begin{bmatrix} P_{k|k}^{T/2} A_k^T & P_{k|k}^{T/2} \\ Q_k^{T/2} & 0 \\ 0 & P_{k+1|T}^{T/2} G_k^T \end{bmatrix} \\ P_{k|T}^{1/2} &= \mathcal{R}_{d+1:2d, d+1:2d} \end{aligned}$$

2.3 Estimation of model parameters

The maximum likelihood estimate of the model parameters $\theta = [q, R, m_{1|0}, P_{1|0}]$ is the maximiser of the likelihood $p(y_{1:T}|\theta)$, or equivalently the minimiser of the ML cost function

$$\phi(\theta) = -\log p(y_{1:T}|\theta).$$

For fixed θ , the cost function can be computed inside the Kalman filter (the first stage of the forward pass of the smoothing algorithm) using

$$\phi(\theta) = \frac{1}{2} \sum_{k=1}^T \sum_{j=1}^{n_k} (\log \det(2\pi S_{k,j}) + v_{k,j}^T S_{k,j}^{-1} v_{k,j}). \quad (4)$$

In the Expectation-Maximisation (EM) method the ML estimate is found by iteratively maximizing a lower bound on the likelihood. An EM method for state-space model parameters that uses a smoother to marginalise the state variables is presented in [14,9]. This method needs to be extended for the ND state space model, which has varying dynamic model matrices and a single-parameter process noise matrix; this is done in the appendix. The EM parameter update formulas are

$$q = \frac{1}{(T-1)d} \sum_{k=1}^{T-1} \text{tr}(\hat{Q}_k \bar{Q}_k^{-1}), \quad R = \frac{1}{N} \sum_{k=1}^T \sum_{j=1}^{n_k} \hat{R}_{k,j}, \quad (5a)$$

$$m_{1|0} = \hat{m}_{1|T}, \quad P_{1|0} = \hat{P}_{1|T}, \quad (5b)$$

where

$$\begin{aligned} \hat{Q}_k &= [I, -A_k] \begin{bmatrix} \hat{m}_{k+1|T} \\ \hat{m}_{k|T} \end{bmatrix} \begin{bmatrix} \hat{m}_{k+1|T} \\ \hat{m}_{k|T} \end{bmatrix}^T \\ &\quad + \begin{bmatrix} \hat{P}_{k+1|T} & \hat{P}_{k+1|T} \hat{G}_k^T \\ \hat{G}_k \hat{P}_{k+1|T} & \hat{P}_{k|T} \end{bmatrix} [I, -A_k]^T, \end{aligned} \quad (6a)$$

$$\hat{R}_{k,j} = (y_{k,j} - H \hat{m}_{k|T})(y_{k,j} - H \hat{m}_{k|T})^T + H \hat{P}_{k|T} H^T, \quad (6b)$$

and the “hat” variables are computed by the smoother with the previous iterand of $\theta = [q, R, m_{1|0}, P_{1|0}]$.

Formulas (6) can also be written in the form

$$\begin{aligned} \hat{Q}_k &= (\hat{m}_{k+1|T} - A_k \hat{m}_{k|T})(\hat{m}_{k+1|T} - A_k \hat{m}_{k|T})^T \\ &\quad + (A_k G_k \hat{Q}_k^{1/2})(A_k G_k \hat{Q}_k^{1/2})^T \\ &\quad + ((I - A_k G_k)(\hat{P}_{k+1|T}^{1/2} + A_k \hat{P}_{k|k}^{1/2})) \\ &\quad \quad ((I - A_k G_k)(\hat{P}_{k+1|T}^{1/2} + A_k \hat{P}_{k|k}^{1/2}))^T, \end{aligned} \quad (7a)$$

$$\begin{aligned} \hat{R}_{k,j} &= (y_{k,j} - H \hat{m}_{k|T})(y_{k,j} - H \hat{m}_{k|T})^T \\ &\quad + (H \hat{P}_{k|T}^{1/2})(H \hat{P}_{k|T}^{1/2})^T. \end{aligned} \quad (7b)$$

Each term in (7) is a product of a matrix with its transpose. Linear algebra software libraries include codes to compute products of this form efficiently and with exact preservation of symmetry; for example the multiplication operator $*$ in MATLAB is overloaded to do this.

2.4 Dense output

The posterior estimate of the state at an inter-abscissa time $t_{k+\theta} = t_k + \theta \Delta_k$, with $0 < \theta < 1$, conditional on the measurements at times up to and including t_k , is denoted

$$x_{k+\theta|k} \sim \mathcal{N}(m_{k+\theta|k}, P_{k+\theta|k}).$$

Its parameters can be obtained using the RTS smoother forward pass formulas by omitting the measurement update stage and applying a dynamic update stage with the modified dynamic model

$$m_{k+\theta|k} = A_{k,\theta} m_{k|k} \quad (8a)$$

$$P_{k+\theta|k} = A_{k,\theta} P_{k|k} A_{k,\theta}^T + Q_{k,\theta} \quad (8b)$$

where $A_{k,\theta} = \exp(F\theta\Delta_k)$ and $Q_{k,\theta} = q\bar{Q}_{k,\theta}$ with

$$\bar{Q}_{k,\theta} = \int_0^{\theta\Delta_k} \exp(F(\theta\Delta_k - \tau)) L L^T \exp(F^T(\theta\Delta_k - \tau)) d\tau.$$

That is, the formulas for the modified model matrices are obtained by using $\theta\Delta_k$ in place of Δ_k in the formulas for

the dynamic transition matrix and process noise covariance given earlier.

The posterior estimate of the interpolatory state conditional on *all* the measurements,

$$x_{k+\theta|T} \sim N(m_{k+\theta|T}, P_{k+\theta|T}),$$

is obtained using the backward pass formula to go from t_{k+1} to $t_{k+\theta}$ instead of to t_k :

$$G_{k,\theta} = P_{k+\theta|k} A_{k,1-\theta}^T P_{k+1|k}^{-1} \quad (9a)$$

$$m_{k+\theta|T} = m_{k+\theta|k} + G_{k,\theta}(m_{k+1|T} - m_{k+1|k}) \quad (9b)$$

The functional form of the interpolant can be inferred from these formulas. Substituting (8) and (9a) into (9b) gives

$$m_{k+\theta|T} = A_{k,\theta} m_{k|k} + (A_{k,\theta} P_{k|k} A_{k,\theta}^T + Q_{k,\theta}) A_{k,1-\theta}^T P_{k+1|k}^{-1} (m_{k+1|T} - m_{k+1|k}).$$

Because the coefficients of $A_{k,\theta}$, $A_{k,1-\theta}$ and $Q_{k,\theta}$ are polynomials in Δ_k , so is the interpolant $m_{k+\theta|T}$. In particular, its first component (the displacement) is a polynomial of degree $2d - 1$.

2.5 Initial parameters

Although EM has good theoretical convergence properties, the convergence can be slow. This slowness can be offset by making a reasonably good choice of initial parameter values. In the MATLAB implementation, the initial iterands for the state $m_{1|0}$ and the measurement noise variance R are set by least-squares fitting a straight line through the first 10 abscissas. The covariance $P_{1|0}$ is set to a tiny multiple of the identity matrix. The driving noise intensity q is then set by minimizing the negative log likelihood, a univariate minimization whose cost function (4) is computed using a Kalman filter.

3 Tests

Corradini et al. [4] compare ND algorithms using five test functions that resemble experimental measurements of different kinds of human movement. They considered different measurement noise levels and sampling rates, and found no large differences in accuracy between the five algorithms that they tested. They however identify two algorithms, which they label F1 and F2, as being the most accurate: the smoothing heptic spline of [16] (widely used because its code is freely available) and the fixed-lag Kalman smoother of [6] with three states. These are also the only algorithms in their tests that are automatic, except that the measurement noise variance needs to be specified.

test	method	displ.	vel.	accel.
T1	F1	0.14	7.27	45.5
	F2	0.13	6.10	36.8
	new	0.15	2.99	11.1
T2	F1	3.51	9.64	25.9
	F2	3.55	9.94	26.7
	new	2.64	8.28	24.8
T3	F1	3.02	9.13	26.1
	F2	3.04	9.40	26.9
	new	2.26	9.37	24.3
T4	F1	2.32	10.40	30.7
	F2	2.39	10.83	34.3
	new	1.77	8.22	33.4
T5	F1	1.95	10.34	39.5
	F2	1.87	9.27	36.0
	new	1.12	6.42	20.6

Table 1
Relative RMS errors (in percentage) of estimates using synthetic displacement data generated from five test functions.

Table 1 shows the errors of displacement, velocity, and acceleration estimates reported in [4] for 94-point noisy displacement sequences generated from five test functions. The error of the estimate of the derivative sequence is reported as the percentage of RMS error relative to the true sequence's RMS value. Also shown are the errors found with the proposed algorithm with $d = 3$ states. The EM iterations were repeated until the norm of the change in the displacement estimate was less than 0.1% of the norm of the estimate; no more than 3 EM iterations were needed in any of the tests.

The methods' errors are not precisely comparable, because different random number generators were used to produce the measurement noise for the data sequences. However, the results indicate that the accuracies of the proposed method are roughly as good and in some cases clearly better than those of the reference methods.

4 Conclusions

The algorithm presented here is based on the integrated Wiener process, which as argued in [7] is a principled and flexible signal model for estimation of derivatives from noisy time series. The proposed ND algorithm has some advantages over that of [7]: it uses a numerically stable square-root smoother algorithm, allows non-equally spaced and simultaneous data, and its implementation is freely available. Also, the ML parameters are computed using a reliable EM iteration, which gives an automatic ND algorithm whose accuracy is as good or better than other methods.

The assumption of additive Gaussian noise may be inadequate for measurements with sporadic outliers. This shortcoming could be addressed by replacing the RTS smoother by a Student-t smoother [12].

References

- [1] K. Ahnert and M. Abel. Numerical differentiation of experimental data: local versus global methods. *Computer Physics Communications*, 177:764–774, 2007.
- [2] Y. Bar-Shalom, X. R. Li, and T. Kirubarajan. *Estimation With Applications to Tracking and Navigation*. Wiley, 2001.
- [3] R. H. Brown, S. C. Schneider, and M. G. Mulligan. Analysis of algorithms for velocity estimation from discrete position versus time data. *IEEE Trans. on Industrial Electronics*, 39(1):11–19, 1992.
- [4] M. L. Corradini, S. Fioretti, and T. Leo. Numerical differentiation in movement analysis: how to standardise the evaluation of techniques. *Medical and Biological Engineering and Computing*, 31:187–197, 1993.
- [5] M. D’Amico and G. Ferrigno. Comparison between the more recent techniques for smoothing and derivative assessment in biomechanics. *Medical and Biological Engineering and Computing*, 30(2):193–204, 1992.
- [6] S. Fioretti and L. Jetto. Accurate derivative estimation from noisy data: a state-space approach. *International Journal of Systems Science*, 20(1):33–53, 1989.
- [7] S. Fioretti and L. Jetto. Low a-priori statistical information model for optimal smoothing and differentiation of noisy signals. *Int. J. of Adaptive Control and Signal Processing*, 8:305–320, 1994.
- [8] G. Giakas and V. Baltzopoulos. A comparison of automatic filtering techniques applied to biomechanical walking data. *J. of Biomechanics*, 30(8):847–850, 1997.
- [9] S. Gibson and B. Ninness. Robust maximum-likelihood estimation of multivariable dynamic systems. *Automatica*, 41, 2005.
- [10] J. Magnus and H. Neudecker. *Matrix Differential Calculus With Applications in Statistics and Econometrics*. Wiley, third edition, 2007.
- [11] J. C. Pezzack, R. W. Norman, and D. A. Winter. An assessment of derivative determining techniques used for motion analysis. *J. of Biomechanics*, 10(5-6):377–382, 1977.
- [12] R. Piché, S. Särkkä, and J. Hartikainen. Robust outlier-robust filtering and smoothing for nonlinear systems using the multivariate student-t distribution. In *IEEE Int. Workshop on Machine Learning for Signal Processing*, 2012.
- [13] L. J. Puglisi, R. J. Saltaren, and C. E. Garcia Cena. On the velocity and acceleration estimation from discrete time-position signal of linear encoders. *J. of Control Engineering and Applied Informatics*, 17(3):30–40, 2015.
- [14] R. H. Shumway and D. S. Stoffer. An approach to time series smoothing and forecasting using the EM algorithm. *J. of Time Series Analysis*, 3(4):253–264, 1982.
- [15] J. A. Walker. Estimating velocities and accelerations of animal locomotion: a simulation experiment comparing numerical differentiation algorithms. *J. of Experimental Biology*, 201(7):981–995, 1998.
- [16] H. J. Woltring. A FORTRAN package for generalized, cross-validatory spline smoothing and differentiation. *Advances in Engineering Software*, 8:104–113, 1986. Source code and Matlab interfaces <http://isbweb.org/software/sigproc.html>.

A Derivation of EM update formulas

Substituting the state space model’s data log-likelihood

$$\log p(x_{1:T}, y_{1:T} | \theta) = \log p(x_1 | \theta) + \sum_{k=1}^{T-1} \log p(x_{k+1} | x_k, \theta) + \sum_{k=1}^T \sum_{j=1}^{n_k} \log p(y_{k,j} | x_k, \theta)$$

into the EM objective function

$$\mathcal{Q}(\theta, \hat{\theta}) = \int p(x_{1:T} | y_{1:T}, \hat{\theta}) \log p(x_{1:T}, y_{1:T} | \theta) dx_{1:T}$$

(where $\hat{\theta}$ is the previous iteration’s parameter value) gives

$$\begin{aligned} \mathcal{Q}(\theta, \hat{\theta}) = & \int p(x_1 | y_{1:T}, \hat{\theta}) \log p(x_1 | \theta) dx_1 \\ & + \sum_{k=1}^{T-1} \int p(x_{k+1}, x_k | y_{1:T}, \hat{\theta}) \log p(x_{k+1} | x_k, \theta) dx_{k+1} \\ & + \sum_{k=1}^T \sum_{j=1}^{n_k} \int p(x_k | y_{1:T}, \hat{\theta}) \log p(y_{k,j} | x_k, \theta) dx_k dx_k. \end{aligned}$$

This is a sum of expectations of log terms. From (3), (1), (2), the log terms are

$$\begin{aligned} \log p(x_1 | \theta) = & -\frac{1}{2} \log \det(2\pi P_{1|0}) \\ & -\frac{1}{2} (x_1 - m_{1|0})^\top P_{1|0}^{-1} (x_1 - m_{1|0}), \\ \log p(x_{k+1} | x_k, \theta) = & -\frac{1}{2} \log \det(2\pi q \bar{Q}_k) \\ & -\frac{1}{2} (x_{k+1} - A_k x_k)^\top (q \bar{Q}_k)^{-1} (x_{k+1} - A_k x_k), \\ \log p(y_{k,j} | x_k, \theta) = & -\frac{1}{2} \log \det(2\pi R) \\ & -\frac{1}{2} (y_{k,j} - H x_k)^\top R^{-1} (y_{k,j} - H x_k). \end{aligned}$$

The distributions with respect to which the expectations are taken are

$$\begin{aligned} x_1 & \sim N(\hat{m}_{1|0}, \hat{P}_{1|0}), \\ \begin{bmatrix} x_{k+1} \\ x_k \end{bmatrix} & \sim N\left(\begin{bmatrix} \hat{m}_{k+1|T} \\ \hat{m}_{k|T} \end{bmatrix}, \begin{bmatrix} \hat{P}_{k+1|T} & \hat{P}_{k+1|T} \hat{G}_k^\top \\ \hat{G}_k \hat{P}_{k+1|T} & \hat{P}_{k|T} \end{bmatrix}\right), \\ x_k & \sim N(\hat{m}_{k|T}, \hat{P}_{k|T}), \end{aligned}$$

where hats indicate values that are computed by the smoothing algorithm applied to the model having parameters $\hat{\theta} = [\hat{q}, \hat{R}, \hat{m}_{1|0}, \hat{P}_{1|0}]$. Computing the expectations gives the for-

mula for the EM objective function as

$$\begin{aligned}\mathcal{Q}(\theta, \hat{\theta}) &= -\frac{1}{2} \log \det(2\pi P_{1|0}) - \frac{1}{2} \text{tr}(P_{1|0}^{-1} \hat{P}_{1|T}) \\ &\quad - \frac{1}{2} (\hat{m}_{1|T} - m_{1|0})^\top P_{1|0}^{-1} (\hat{m}_{1|T} - m_{1|0}) \\ &\quad - \frac{1}{2} \sum_{k=1}^{T-1} \left(\log \det(2\pi Q_k) + \text{tr}(Q_k^{-1} \hat{Q}_k) \right) \\ &\quad - \frac{1}{2} \sum_{k=1}^T \sum_{j=1}^{n_k} \left(\log \det(2\pi R) + \text{tr}(R^{-1} \hat{R}_{k,j}) \right),\end{aligned}$$

where \hat{Q}_k and $\hat{R}_{k,j}$ are given by (6). Using standard matrix differential calculus formulas [10], the partial derivatives of the EM objective function are

$$\begin{aligned}\partial \mathcal{Q}(\theta, \hat{\theta}) / \partial q &= \frac{1}{2} \sum_{k=1}^{T-1} \text{tr} \left(Q_k^{-1} \frac{\partial Q_k}{\partial q} (-I + \hat{Q}_k Q_k^{-1}) \right) \\ &= \frac{1}{2} \left(-\frac{(T-1)d}{q} + \frac{1}{q^2} \sum_{k=1}^{T-1} \text{tr}(\hat{Q}_k \bar{Q}_k^{-1}) \right), \\ \partial \mathcal{Q}(\theta, \hat{\theta}) / \partial R &= \frac{1}{2} \sum_{k=1}^T \sum_{j=1}^{n_k} R^{-1} (-I + \hat{R}_{k,j} R^{-1}), \\ \partial \mathcal{Q}(\theta, \hat{\theta}) / \partial m_{1|0}^\top &= (\hat{m}_{1|T} - m_{1|0})^\top P_{1|0}^{-1}, \\ \partial \mathcal{Q}(\theta, \hat{\theta}) / \partial P_{1|0} &= \frac{1}{2} P_{1|0}^{-1} \left(-I + (\hat{P}_{1|T} \right. \\ &\quad \left. + (\hat{m}_{1|T} - m_{1|0})(\hat{m}_{1|T} - m_{1|0})^\top) \right) P_{1|0}^{-1}.\end{aligned}$$

Setting these to zero and solving gives the EM update formulas (5–6).

The covariance matrix in (6a) can be written as

$$\begin{aligned}&\begin{bmatrix} \hat{P}_{k+1|T} & \hat{P}_{k+1|T} \hat{G}_k^\top \\ \hat{G}_k \hat{P}_{k+1|T} & \hat{P}_{k|T} \end{bmatrix} \\ &= \begin{bmatrix} I & 0 \\ \hat{G}_k & I \end{bmatrix} \begin{bmatrix} \hat{P}_{k+1|T} & 0 \\ 0 & \hat{P}_{k|T} - \hat{G}_k \hat{P}_{k+1|T} \hat{G}_k^\top \end{bmatrix} \begin{bmatrix} I & 0 \\ \hat{G}_k & I \end{bmatrix}^\top.\end{aligned}$$

Substituting the identities

$$\hat{P}_{k|T} = \hat{P}_{k|k} + \hat{G}_k (\hat{P}_{k+1|T} - \hat{P}_{k+1|k}) \hat{G}_k^\top$$

and

$$\hat{P}_{k+1|k} = A_k \hat{P}_{k|k} A_k^\top + \hat{Q}_k,$$

and applying the Joseph formula, the last element of the diagonal matrix can be rewritten as

$$\begin{aligned}\hat{P}_{k|T} - \hat{G}_k \hat{P}_{k+1|T} \hat{G}_k^\top &= \hat{P}_{k|k} - \hat{G}_k \hat{P}_{k+1|k} \hat{G}_k^\top \\ &= \hat{P}_{k|k} - \hat{G}_k (A_k \hat{P}_{k|k} A_k^\top + \hat{Q}_k) \hat{G}_k^\top \\ &= (I - \hat{G}_k A_k) \hat{P}_{k|k} (I - \hat{G}_k A_k)^\top + \hat{G}_k \hat{Q}_k \hat{G}_k^\top.\end{aligned}$$

Formula (7a) is then obtained by replacing the covariance matrices by their Cholesky factorisations.